



(12) **United States Patent**
LaForest

(10) **Patent No.:** **US 9,268,608 B2**
(45) **Date of Patent:** **Feb. 23, 2016**

(54) **AUTOMATIC ADMINISTRATION OF UNIX COMMANDS**

(75) Inventor: **Mark LaForest**, Austin, TX (US)

(73) Assignee: **Oracle International Corporation**,
Redwood Shores, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 680 days.

(21) Appl. No.: **12/393,894**

(22) Filed: **Feb. 26, 2009**

(65) **Prior Publication Data**

US 2010/0217849 A1 Aug. 26, 2010

(51) **Int. Cl.**
G06F 15/173 (2006.01)
G06F 9/50 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 9/5016** (2013.01)

(58) **Field of Classification Search**
CPC . G06F 9/5016; G06F 9/30101; G06F 9/3836;
G06F 9/3838; G06F 9/3851; G06F 15/8007;
G06Q 10/063; H04L 29/08171; Y10S
707/99945
USPC 709/201, 202, 223, 226, 238; 719/313,
719/314
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,297,262 A * 3/1994 Cox et al. 710/36
5,442,791 A 8/1995 Wrabetz et al.
5,860,143 A * 1/1999 Plakosh et al. 711/206
6,782,537 B1 * 8/2004 Blackmore G06F 9/524
718/106
7,428,239 B1 * 9/2008 Seigneurbieux H04L 12/5693
370/418

7,627,736 B2 12/2009 Stuttard et al. 712/10
7,739,366 B2 * 6/2010 Ansari 709/223
7,739,432 B1 * 6/2010 Shaw et al. 710/74
8,266,633 B1 * 9/2012 Saulsbury G06F 9/45558
718/1
2002/0156917 A1 10/2002 Nye
2002/0198921 A1 12/2002 Jameson
2003/0120811 A1 6/2003 Hanson et al.
2003/0145050 A1 7/2003 Block et al.
2003/0233571 A1 * 12/2003 Kraus et al. 713/200
2004/0215398 A1 * 10/2004 Mixon et al. 702/19
2006/0017969 A1 * 1/2006 Ly et al. 358/1.15
2006/0212422 A1 9/2006 Khilani et al.
2006/0265710 A1 * 11/2006 Ansari 718/1

(Continued)

OTHER PUBLICATIONS

Global, Internet Archive Wayback Machine, retrieved from <http://web.archive.org/web/20031209063419/http://users.sitestar.net/~efinch/global.html> on Jan. 11, 2011, 2 pages.

(Continued)

Primary Examiner — Phuoc Nguyen

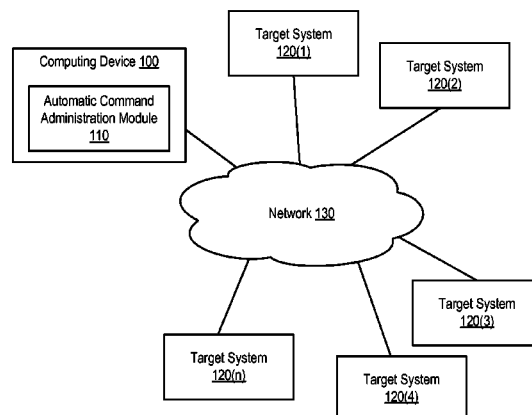
Assistant Examiner — Davoud Zand

(74) *Attorney, Agent, or Firm* — Kilpatrick Townsend & Stockton LLP

(57) **ABSTRACT**

Various techniques for automatically administering UNIX commands to target systems are disclosed. One method involves receiving information identifying a UNIX command and additional information identifying one or more target systems. The method then issues N instances of the UNIX command in parallel to the one or more target systems, where N is an integer greater than one. The N instances of the UNIX command are issued automatically, in response to receipt of the information and the additional information. In some situations, issuing the N instances of the UNIX command in parallel involves creating N threads, where each of the N threads is configured to issue a respective one of the N instances of the UNIX command to a respective one of the target systems.

22 Claims, 3 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

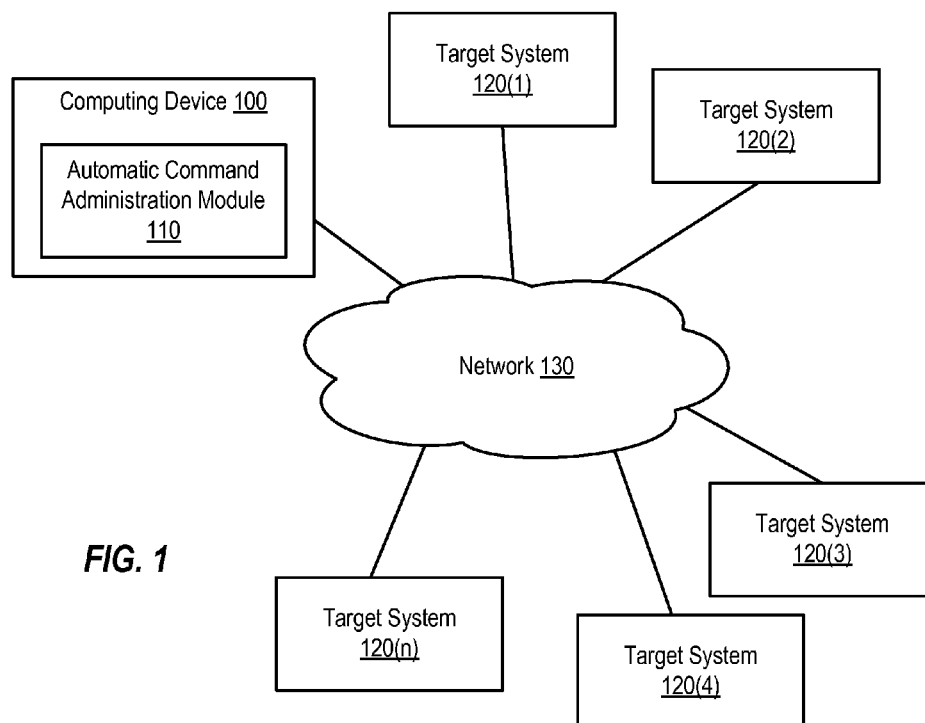
2007/0118601 A1* 5/2007 Pacheco H04L 12/5693
709/206
2007/0130303 A1* 6/2007 Anna et al. 709/223
2007/0255829 A1* 11/2007 Pecus et al. 709/225
2008/0040564 A1* 2/2008 Kubo et al. 711/154
2008/0235686 A1* 9/2008 Brenner 718/101
2009/0019538 A1* 1/2009 Pandya 726/13
2009/0055603 A1* 2/2009 Holt G06F 15/16
711/154
2010/0011372 A1* 1/2010 Trench 718/107
2010/0100624 A1* 4/2010 Somogyi et al. 709/226
2010/0154042 A1* 6/2010 Miyamoto et al. 726/7
2010/0199257 A1* 8/2010 Biggerstaff G06F 8/76

2010/0217849 A1* 8/2010 LaForest 717/104
2011/0078108 A1 3/2011 Kumar 709/222
2011/0099146 A1 4/2011 Mcalister et al.

OTHER PUBLICATIONS

Global, v1.03—Ed Finch (efinch@newsreal.com), Jun. 1999, 12 pages.
U.S. Appl. No. 13/306,256, filed Nov. 29, 2011, Salas.
U.S. Appl. No. 13/306,256, “Final Office Action”, Jan. 15, 2015, 11 pages.
U.S. Appl. No. 13/306,256, “Non Final Office Action”, May 14, 2014, 11 pages.

* cited by examiner



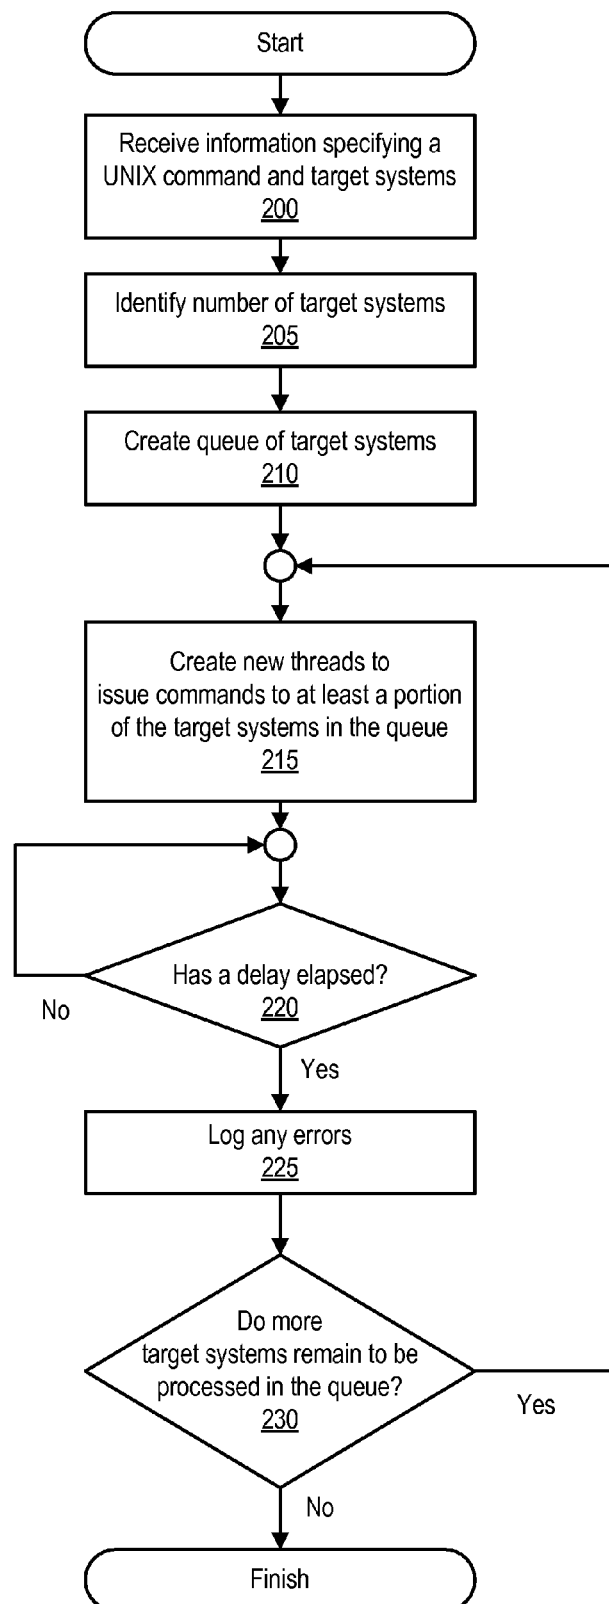
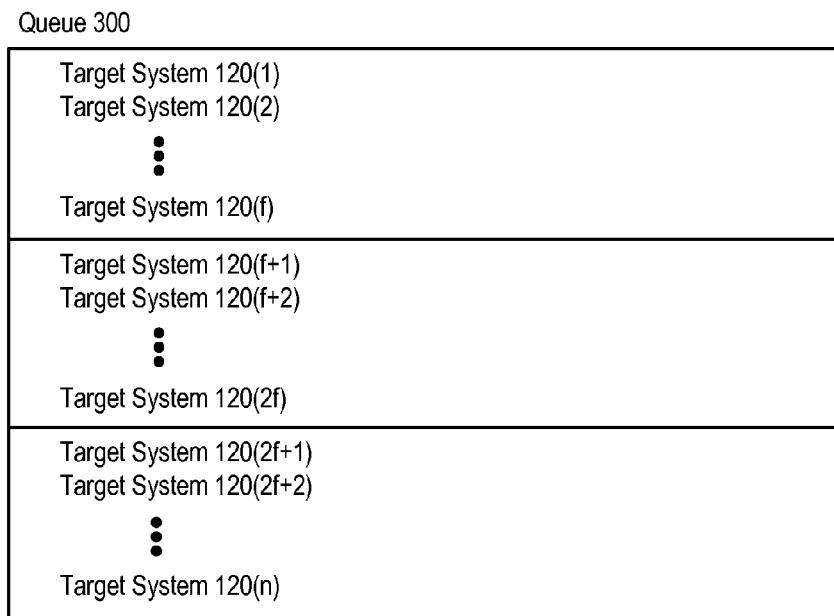
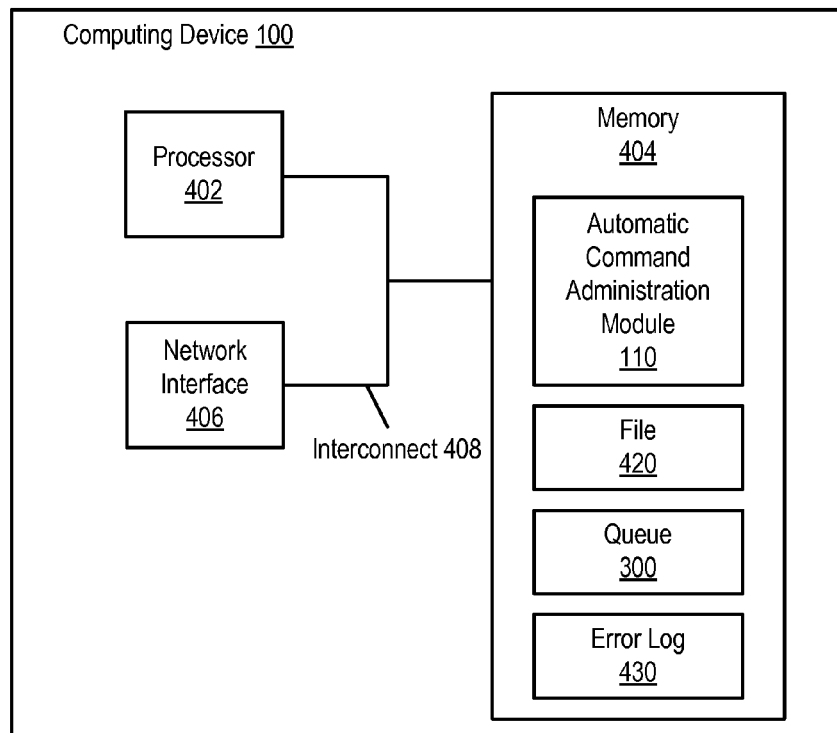


FIG. 2

**FIG. 3****FIG. 4**

1

AUTOMATIC ADMINISTRATION OF UNIX COMMANDS

FIELD OF THE INVENTION

This invention relates to system administration and, more particularly, to performing administrative tasks that involve issuing UNIX commands.

DESCRIPTION OF THE RELATED ART

System administrators are often responsible for issuing UNIX commands to hundreds or even thousands of systems at once. Unfortunately, most of the tools that are currently available to assist with such tasks issue commands serially. As a result, it can take hours or even days for such tools to issue the commands to all of the target systems. Accordingly, new tools that allow system administrators to more quickly issue commands to target systems are desirable.

SUMMARY OF THE INVENTION

Various systems and methods for automatically issuing UNIX commands to target systems are disclosed. In one embodiment, a method involves receiving information identifying a UNIX command and additional information identifying one or more target systems. The method then issues N instances of the UNIX command in parallel to the one or more target systems, in response to receiving the information and the additional information. The commands can be issued via a network.

The operations of receiving the information and issuing the N instances of the UNIX command can be performed by a computing device, in response to the computing device executing a compiled application that implements an automatic command administration module. N is an integer greater than one. N can be significantly greater than one in some situations. For example, N can be greater than one hundred or even greater than one thousand.

In some embodiments, issuing the N instances of the UNIX command in parallel involves creating N threads. Each of the N threads is configured to issue a respective one of the N instances of the UNIX command. Before creating the N threads, an amount of available memory can be determined. The amount of available memory can be compared to a required amount of memory needed to create the N threads. Creation of the N threads can be conditioned on whether the amount of available memory exceeds the required amount of memory needed to create the N threads.

The method can additionally involve adding information identifying one of the target systems to an error log, in response to detecting that a corresponding one of the N threads has yet to complete. The detection of whether the corresponding thread has completed can be performed in response to expiration of a delay period, which begins when the N threads are created. Alternatively (or additionally), error information identifying one of the target systems can be logged in response to a thread receiving an error indication from that target system.

In some embodiments, the additional information identifying the target systems is stored in a file. Receiving the additional information can thus involve receiving information, such as a file name, identifying the file and then accessing the file.

The UNIX command can be a multipart command. The information specifying the UNIX command can identify the order of each of several parts of the multipart command.

2

Issuing the N instances of the UNIX command can involve issuing each of the parts of the multipart command in the order identified by the information.

In one embodiment, the method involves creating a queue of target systems. Issuing the N instances of the UNIX command involves issuing an instance of the UNIX command to each of the first N target systems in the queue. The value of N can, in some embodiments, be specified as part of the information.

If the queue contains more than N target systems, the method can also involve issuing the instance of the UNIX command to each of a second N target systems in the queue. The instance of the UNIX command is issued to the first N target systems at a first time, and the instance of the UNIX command is issued to the second N target systems at a second time. The second time is delayed relative to the first time. In some embodiments, the information can indicate the amount of the delay.

The foregoing is a summary and thus contains, by necessity, simplifications, generalizations and omissions of detail; consequently those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way limiting. Other aspects, inventive features, and advantages of the present invention, as defined solely by the claims, will become apparent in the non-limiting detailed description set forth below.

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of the present invention may be acquired by referring to the following detailed description and the accompanying drawings, listed below, in which like reference numbers indicate like features.

FIG. 1 is a block diagram of a computer system that can remotely administer multiple target systems, according to one embodiment of the present invention.

FIG. 2 is a flow chart of one embodiment of a method of remotely administering multiple target systems.

FIG. 3 is an example of a queue that can be generated during the remote administration of multiple target systems, according to one embodiment of the present invention.

FIG. 4 is a block diagram of a computing device, which illustrates how an automatic command administration module can be implemented in software, according to one embodiment of the present invention.

While the invention is susceptible to various modifications and alternative forms, specific embodiments of the invention are provided as examples in the drawings and detailed description. It should be understood that the drawings and detailed description are not intended to limit the invention to the particular form disclosed. Instead, the intention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the invention as defined by the appended claims.

DETAILED DESCRIPTION

FIG. 1 is a block diagram of a computer system that can remotely administer multiple target systems. As shown, the system includes a computing device 100 that implements an automatic command administration module 110. Computing device 100 is coupled to target systems 120(1), 120(2), 120(3), 120(4), and 120(n) by network 130. In this example, "n" is an integer that is greater than or equal to five (5), indicating that the system can contain five or more target systems. In general, the number of target systems can vary

among embodiments. For example, some systems may contain relatively few target systems, while others may include thousands of target systems.

Each target system **120(1)-120(n)** is a computing device. Computing devices can include personal computers, network appliances, servers, personal digital assistants, mobile phones, storage controllers, laptop computers, and the like. In many embodiments, most or all of the target system computing devices will be servers (e.g., in a data center or other environment that requires a significant number of servers to operate).

Some target systems can be virtual target systems. These systems behave as if they are independent physical machines and have their own independent identities (e.g., in the form of hostnames, network addresses, and the like). However, these systems may be included in the same physical chassis as one or more other target systems, or even in the same physical chassis as computing device **100**.

Automatic command administration module **110** is configured to receive user input specifying a UNIX command, as well as input specifying a set of target systems. Based upon this input, automatic command administration module **110** is configured to automatically, without further user interaction, administer the identified target systems by issuing instances of the UNIX command, in parallel, to at least some of the target systems. Automatic command administration module **110** can continue to issue instances of the UNIX command, in parallel, to target systems until all target systems have been issued an instance of the UNIX command.

As used herein, a UNIX command is any command that is designed to be performed by a UNIX or UNIX-like operating systems. Generally, UNIX operating systems are proprietary operating systems, such as Solaris, Hewlett Packard UNIX (HP-UX), and Advanced Interactive eXecutive (AIX), that fully comply with the family of standards known as the Single UNIX Specification (SUS), the core of which is maintained by the Austin Common Standards Revision Group (CSRG). UNIX-like operating systems are typically non-proprietary operating systems such as Linux and BSD (Berkeley Software Distribution, or Berkeley Unix).

Automatic command administration module **110** is implemented as a compiled software application. Since automatic command administration module **110** is already compiled (in contrast with a script), automatic command administration module **110** does not need to be compiled at run time before being used. This in turn reduces the amount of system resources required to operate automatic command administration module **110**. Furthermore, as a compiled application (in contrast with a script), automatic command administration module **110** can control the number of instances of the command that are issued in parallel (e.g., by controlling the number of threads and/or forks, as discussed in more detail below), based upon conditions such as current memory usage.

In some embodiments, automatic command administration module **110** is configured to issue the commands to multiple target systems in parallel by creating multiple threads, each of which is configured to issue a command to a respective target system. These threads can be created using a command, such as “pthread_create”, that adds a new thread to the current process. When this thread creation technique is used, multiple threads can coexist within the same process. Alternatively, each of the threads can be created by issuing a fork command, which creates a new child process. This technique causes each thread to exist in its own child process. It is noted that in some systems, the operating system may convert thread commands into fork commands.

In some embodiments, automatic command administration module **110** is invoked by a command entered via a command line interface (CLI) of an operating system (alternatively, the automatic command administration module **110** can be invoked via a graphical user interface (GUI), after which a user is prompted for the command to issue, the target systems to which the command should be issued, and other desired parameters). This command can include a command identifier, which identifies automatic command administration module **110** (e.g., by identifying the file name of the executable file that is executed to implement automatic command administration module **110**). This command can also take one or more arguments or parameters. Such arguments can include the name of the command to be provided to the target systems, the identities of the target systems (or where to find these identities), the maximum number of commands to issue in parallel, and the like.

Some arguments can be specified using command line switches. Command line switches are special arguments, each of which has an identifier (typically a single letter preceded by a minus sign, such as -f) and, optionally, a user-specified value. Command line switches are typically placed prior to any other arguments. In one embodiment, command line switches are available to specify the maximum number of commands to issue in parallel, the part(s) of the command that is to be issued to the target systems, the delay between issuing successive parts of a command to a target system and/or between issuing commands and then determining whether the issued commands were successful, the name of a file that stores information (e.g., hostnames) identifying the target systems, and whether platform-specific arguments are to be included with the command issued to a given target system. A command line switch (e.g., -h) may also be available to request help (e.g., in the form of a list of the available command line switches for the automatic command administration module and the allowable values that can be provided with each). Some command line switches, such as a switch used to specify part of a multipart command, should be accompanied by a value, such as the name of the part of the command. Other command line switches, such as a switch used to request help, should not be accompanied by any values.

In some embodiments, automatic command administration module **110** supports multipart commands. In such embodiments, when a user invokes automatic command administration module **110**, the user can specify each part of the multipart command, as well as the order in which the parts are to be performed. In one embodiment, this is facilitated through the use of numbered command line switches (e.g., numbered 1, 2, 3, . . . , n, where n is the number of parts in the multipart command). For example, to cause automatic command administration module **110** to issue a two-part command, the user can specify “-1 part1 -2 part2”. This indicates that command “part1” should be sent to a target system prior to command “part2.” As noted above, another command line switch can be used to specify the delay between when a part of a command is issued and when a determination is made as to whether that part of the command was issued successfully. This delay can also set a minimum amount of time that can elapse between when successive parts of a command are issued.

As noted above, instead of being provided to automatic command administration module **110** directly, the identities of the target systems to be administered can be stored in a file, and the name of this file can then be provided to automatic command administration module **110** (e.g., by specifying the filename and/or pathname in an appropriate command line

5

switch). This technique is useful when a large number of (e.g., more than ten) target systems are present. The file can store the identity of each target system, as well as any platform-specific parameters to be included with the command sent to each target system. The parameters, if any, can vary for each identified target system. Thus, the file can specify that Hostname1 is associated with ParameterA, that Hostname2 is associated with ParameterB and ParameterC, that Hostname3 is associated with ParameterA, and so on.

Automatic command administration module 110 can parse this file in order to obtain the identity of each target system to which a command should be issued, as well as to obtain any platform-specific parameters to be provided to each target system. The platform-specific parameters are platform-specific in that these parameters are specific to the operating system of the target system to which the parameters should be sent.

FIG. 2 provides more details regarding how an automatic command administration module, like automatic command administration module 110 of FIG. 1, can issue commands to target systems. Operations 200-230 of FIG. 2 can be performed by an automatic command administration module like that shown in FIG. 1.

The method begins at operation 200, when information specifying a UNIX command and two or more target systems is received. As noted above, this information can be received via a user interface. If the user interface is a CLI, the information can be provided in the form of a command that invokes an automatic command administration module and a list of parameters. The information specifying the UNIX command to issue to the target systems can be provided directly as one of the parameters.

In some situations, the information identifying the target systems can be provided indirectly. For example, a parameter of the command invoking the automatic command administration module can specify a file, which can in turn store the identities of the target systems. Thus, receiving the information identifying the target systems can involve first receiving information identifying a file, and then receiving (e.g., by reading or otherwise accessing) the identified file.

It is noted that additional information can also be received along with the command to be issued to the target systems (e.g., such information can be received as one or more parameters of the command used to invoke the automatic command administration module). Such information can specify the different parts of the command, and the order in which those parts should be issued, if the command is a multipart command. Similarly, such information can include information specifying a delay (e.g., the amount of time to allow to elapse before checking to see if the command was issued successfully and/or the amount of time to allow to elapse between issuing successive parts of a command). Other information can include information indicating whether platform-specific parameters should be issued to the target systems, information identifying the maximum number of commands to be issued in parallel, and the like.

At 205, the number of identified target systems (as identified in the information received at 200) is determined. Determining this number can involve counting the number of target systems specified as parameters of the command, if the target systems are specified directly. If the target systems are specified indirectly, determining this number can involve accessing a file, and then processing that file to determine the number of target systems identified within the file.

At 210, a queue of target systems is created. An example of this queue is shown in FIG. 3. The queue lists all of the target systems identified in the information received at 200. For

6

each target system identified at 200, one entry is created in the queue. This entry identifies the target system, the command to be sent to that target system (including all parts of the command, in order, if the command is a multipart command), and any parameters to be provided to that target system with the command.

It is noted that in some situations, the information received at 200 can specify the same target system multiple times (e.g., in situations in which it is desirable to repeatedly perform the same command on the same target system more than once). Accordingly, when a target system is identified multiple times, multiple entries will be created for that target system within the queue created at 210.

At 215, multiple new threads are created. Each thread corresponds to one entry in the queue. Each thread is configured to send the command specified in the thread's corresponding queue entry to the target system specified in the thread's corresponding queue entry. If the command is a multipart command, the thread is configured to issue the parts of the command to the target system, in the order those parts are listed in the corresponding queue entry (and with an appropriate amount of delay between successively issued command parts, if a delay has been specified). Similarly, if the queue entry indicates that additional parameters should be provided to the target system with a particular command or part of a command, the thread will provide those parameters to the target system. If the target system is located remotely from the system that implements the command administration module (e.g., if the target system is not a virtual system implemented within the same physical chassis as the computing device implementing the automatic command administration module), the thread will issue the command to the target system via a network.

As noted in the description of operation 215 of FIG. 2, the new threads may only be able to issue commands to some (as opposed to all) of the target systems listed in the queue. This situation arises if the number of entries in the queue exceeds the number of threads that the automatic command administration module is allowed to create at any point in time. The number of threads that the automatic command administration module is allowed to create can be determined in a variety of ways, based upon one or more of a variety of different factors.

In one embodiment, before creating new threads, the automatic command administration module is configured to check the amount of memory that is available to the automatic command administration module, and to calculate the amount of memory needed to create a desired number of threads. If the amount of memory is insufficient, the automatic command administration module chooses not to create the new threads at that time. This can lead to the automatic command administration module either returning an error indication to the user that invoked the automatic command administration module or delaying performing of the requested administrative task. Alternatively, the automatic command administration module can determine the number (which is less than the desired number) of new threads to create, based upon the amount of available memory.

As noted above, in some situations, a user may specify the maximum number of threads to create at any given time. In this situation, the desired number of threads is the user-specified number. If no user-specified number is provided (e.g., in situations in which the user failed to provide one, or in systems in which users are not allowed to specify maximum numbers of threads), the automatic command administration

module may use a default maximum number that has been preconfigured in the automatic command administration module.

If the number of target systems exceeds the number of threads that can be created at a given time, some target systems will be issued commands later than others. Thus, if there are 10,000 target systems, but only 1,000 threads can be created at the same time, target systems **1-1000** will be issued commands by 1,000 respective threads created during a first period. After those threads finish executing, another 1,000 threads can be created to issue commands to target systems **1001-2000**, and so on.

Each time that a new group of threads is created, the automatic command administration module can allow a delay period (e.g., as specified as a default value, or as specified as a parameter provided to the automatic command administration module) to elapse. The delay period begins when the new threads are created. At the end of this period (as determined at **220**), the automatic command administration module can check to see whether all of the threads have completed successfully, which in turn indicates that the commands issued by those threads were issued successfully. If any errors have occurred (e.g., if a thread has failed to complete by the time the delay period expires, if the thread received an error indication from the target system), the automatic command administration module can log those errors, as shown at **225**. If any threads have failed to complete, the automatic command administration module can also force those threads to cease execution (alternatively, this action can be performed after the entire queue has been processed).

After a particular set of threads has completed execution, the automatic command administration module determines whether additional target systems remain to be processed in the queue. If more target systems remain to be processed within the queue, the method returns to operation **215**. If all of the target systems have been processed, the method is finished.

FIG. 3 is an example of a queue that can be generated during the remote administration of multiple target systems. As shown, the queue lists n target systems **120(1)-120(n)**. In this example, the maximum number of threads that can be created at a given point in time is f . As such, the queue has been subdivided into groups of f target systems. The first group includes target systems **120(1)-120(f)**, the second group includes target systems **120(f +1)-120($2f$)**, and the third group includes target systems **120($2f$ +1)-120(n)** (n may be less than $3f$, and thus there may be fewer than f target systems in the final group within the queue).

Each group of target systems within the queue represents a group of target systems that can be administered in parallel by the automatic command administration module. Thus, the automatic command administration module can create a first group of f threads to administer the first group of target systems **120(1)-120(f)**. After these f threads complete, the automatic command administration module can create a second group of f threads to administer the second group of target systems, and so on.

FIG. 4 is a block diagram of a computing device, which illustrates how an automatic command administration module can be implemented in software. As illustrated, computing device **100** (e.g., as also shown in FIG. 1) includes one or more processors **402** (e.g., microprocessors, programmable logic devices (PLDs), or application specific integrated circuits (ASICs)) configured to execute program instructions stored in memory **404**. Memory **404** can include various types of RAM (Random Access Memory), Read Only Memory (ROM), Flash memory, micro electromechanical systems

(MEMS) memory, magnetic core memory, and the like. Memory **404** can include both volatile and non-volatile memory. Computing device **100** also includes one or more interfaces **406**. Processor **402**, interface **406**, and memory **404** are coupled to send and receive data and control signals by a bus or other interconnect.

Interface **406** can include a network interface to various networks and/or interfaces to various peripheral buses. For example, interface **406** can include a network interface via which a thread created by automatic command administration module **110** can send a command and/or a command parameter to a target system. The thread can also receive responsive indications from the target system (e.g., indicating whether the target system successfully received and/or performed the command sent by the thread). Interface **406** can also include an interface to one or more storage devices (e.g., on which a file, queue, and/or error log are stored).

In this example, program instructions and data executable to implement all or part of automatic command administration module **110** are stored in memory **404**. Memory **404** also stores all or part of a file **420**, a queue **300** (e.g., as shown in FIG. 3), and an error log **430**. Automatic command administration module **110** can access file **420** (e.g., in response to the filename and/or pathname of file **420** being provided as a parameter to command administration module **110**) in order to obtain a list of target systems to which commands should be issued. Automatic command administration module **110** can process file **420** to generate queue **300**, which can then be processed by threads that are configured to issue commands to target systems. If any errors are detected (e.g., because a thread receives an error indication from a target system, or because a thread fails to complete execution within a delay period), automatic command administration module **110** can write information identifying those errors (e.g., such information can identify a target system from which an error indication was received, a target system that was being administered by a failed thread, a timestamp indicating approximately when the error was detected, the command that was being issued to the target system as part of the administrative action, and the like) to error log **430**.

The program instructions and data implementing automatic command administration module **110** can be stored on various computer readable media such as memory **404**. In some embodiments, such software is stored on a computer readable medium such as a compact disc (CD), digital versatile disc (DVD), hard disk, optical disk, tape device, floppy disk, and the like). In order to be executed by processor **402**, the instructions and data can be loaded into memory **404** from the other computer readable medium. The instructions and/or data can also be transferred to computing device **100** for storage in memory **404** via a network such as the Internet or upon a carrier medium.

Although the present invention has been described in connection with several embodiments, the invention is not intended to be limited to the specific forms set forth herein. On the contrary, it is intended to cover such alternatives, modifications, and equivalents as can be reasonably included within the scope of the invention as defined by the appended claims.

What is claimed is:

1. A method comprising:

receiving, by a computing device, input from a user, the input comprising information identifying a multipart command, identifying an order of each part of a plurality of parts of the multipart command, identifying a delay that sets a minimum amount of time to elapse between successive parts of the multipart command being issued,

9

and identifying a maximum number of commands to issue in parallel, the input further comprising additional information identifying N target systems, wherein the multipart command is configured to be performed by an operating system; and

administering, by the computing device, the identified N target systems using the received input, including:

- creating a queue relating to the N target systems, the queue comprising N entries, wherein each of the N entries identifies one of the N target systems, the multipart command to be issued to at least one target system of the N target systems, and a set of parameters to be provided with the multipart command, the set of parameters including platform-specific parameters of each of the N target systems;
- creating N threads, each of the N threads corresponding with one of the N entries in the queue, and each of the N threads being configured to issue a respective one of N instances of the multipart command to a respective one of the N target systems; and
- issuing, using the N threads, each of the parts of the N instances of the multipart command to each of the maximum number of N target systems in parallel, in the order identified by the information, and based on the minimum amount of time associated with the delay, respectively, wherein N is an integer greater than one.

2. The method of claim 1, further comprising:

- determining an amount of memory available; and
- comparing the amount of memory available to a required amount of memory needed to create the N threads, wherein the creating the N threads is performed in response to the amount of memory available exceeding the required amount of memory needed to create the N threads.

3. The method of claim 1, further comprising:

- adding information identifying one of the N target systems to an error log, in response to detecting that a corresponding one of the N threads has yet to complete, wherein
- the detecting is performed in response to expiration of the delay, and
- the delay begins when the N threads are created.

4. The method of claim 1, wherein the additional information is stored in a file, and the receiving the additional information comprises receiving a file name identifying the file and accessing the file.

5. The method of claim 1, wherein N is greater than one hundred.

6. The method of claim 1, wherein the N instances of the multipart command are issued via a network.

7. The method of claim 1, further comprising:

- receiving an error indication from one of the N target systems; and
- adding information identifying the one of the N target systems to an error log.

8. The method of claim 1, wherein issuing the N instances of the multipart command comprises issuing an instance of the multipart command to a first N target systems in the queue.

9. The method of claim 8, further comprising:

- issuing the instance of the multipart command to a second N target systems in the queue, wherein the instance of the multipart command is issued to the first N target systems at a first time, the instance of the multipart command is issued to the second N target systems at a

10

second time, the second time is delayed relative to the first time, and the information indicates an amount of the delay.

10. A non-transitory computer readable storage medium comprising a compiled application, wherein the compiled application is executable by a processor to:

- receive input from a user, the input comprising information identifying a multipart command, identifying an order of each part of a plurality of parts of the multipart command, identifying a delay that sets a minimum amount of time to elapse between successive parts of the multipart command being issued, and identifying a maximum number of commands to issue in parallel, the input further comprising additional information identifying N target systems, wherein the multipart command is configured to be performed by an operating system; and
- administer the identified N target systems using the received input, including:

- creating a queue relating to the N target systems, the queue comprising N entries, wherein each of the N entries identifies one of the N target systems, the multipart command to be issued to at least one target system of the N target systems, and a set of parameters to be provided with the multipart command, the set of parameters including platform-specific parameters of each of the N target systems;
- creating N threads, each of the N threads corresponding with one of the N entries in the queue, and each of the N threads being configured to issue a respective one of N instances of the multipart command to a respective one of the N target systems; and
- issuing, using the N threads, each of the parts of the N instances of the multipart command to each of the maximum number of N target systems in parallel, in the order identified by the information, and based on the minimum amount of time associated with the delay, respectively, wherein N is an integer greater than one.

11. The non-transitory computer readable storage medium of claim 10, wherein the compiled application is further executable to:

- add information identifying one of the N target systems to an error log, in response to detecting that a corresponding one of the N threads has yet to complete, wherein
- the detecting is performed in response to expiration of the delay, and
- delay begins when the N threads are created.

12. The non-transitory computer readable storage medium of claim 10, wherein the compiled application is further executable to:

- determine an amount of memory available; and
- compare the amount of memory available to a required amount of memory needed to create the N threads, wherein the N threads are created in response to the amount of memory available exceeding the required amount of memory needed.

13. The non-transitory computer readable storage medium of claim 10, wherein the additional information is stored in a file, and the additional information is received by receiving a file name identifying the file and accessing the file.

14. The non-transitory computer readable storage medium of claim 10, wherein N is greater than one hundred.

15. The non-transitory computer readable storage medium of claim 10, wherein the N instances of the multipart command are issued via a network.

11

16. The non-transitory computer readable storage medium of claim 10, wherein the compiled application is further executable to:

- receive an error indication from one of the N target systems; and
- add information identifying the one of the N target systems to an error log.

17. The non-transitory computer readable storage medium of claim 10, wherein the N instances of the multipart command are issued by issuing an instance of the multipart command to a first N target systems in the queue.

18. The non-transitory computer readable storage medium of claim 10, wherein the compiled application is further executable to:

- issue the instance of the multipart command to a second N target systems in the queue, wherein the instance of the multipart command is issued to the first N target systems at a first time, the instance of the multipart command is issued to the second N target systems at a second time, the second time is delayed relative to the first time, and the information indicates an amount of delay.

19. A system comprising:

a first circuit for storing a list of N target systems;

a second circuit for receiving input from a user, the input comprising information identifying a multipart command, identifying an order of each part of a plurality of parts of the multipart command, identifying a delay that sets a minimum amount of time to elapse between successive parts of the multipart command being issued, and identifying a maximum number of commands to issue in parallel, the input further comprising additional information identifying the list, wherein the multipart command is configured to be performed by an operating system;

12

a third circuit for creating a queue relating to the N target systems, the queue comprising N entries, wherein each of the N entries identifies one of the N target systems, the multipart command to be issued to at least one target system of the N target systems, and a set of parameters to be provided with the multipart command, the set of parameters including platform-specific parameters of each of the N target systems;

a fourth circuit for creating N threads, each of the N threads corresponding with one of the N entries in the queue, and each of the N threads being configured to issue a respective one of N instances of the multipart command to a respective one of the N target systems; and

a fifth circuit for issuing, using the N threads, each of the parts of the N instances of the multipart command to each of the maximum number of N target systems in parallel, in the order identified by the information, and based on the minimum amount of time associated with the delay, respectively, wherein N is an integer greater than one.

20. The method of claim 1, further comprising:

subdividing the queue into groups of M target systems, wherein

the M target systems are a subset of the N target systems, and

the M target systems are grouped according to the maximum number of threads that can be created at a time.

21. The method of claim 1, wherein the N threads are created by using a second multipart command that adds a new thread to a current process.

22. The method of claim 1, wherein the N threads are created by issuing a second multipart command that creates a new child process.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 9,268,608 B2
APPLICATION NO. : 12/393894
DATED : February 23, 2016
INVENTOR(S) : LaForest


Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In The Claims

In column 10, line 49, in claim 11, before “delay begins” insert -- the --.

Signed and Sealed this
Fifteenth Day of November, 2016

A handwritten signature in black ink, reading "Michelle K. Lee". The signature is fluid and cursive, with the first letters of each name being capitalized and prominent.

Michelle K. Lee
Director of the United States Patent and Trademark Office